ACCEPTED FROM OPEN CALL

STRATEGIES FOR ADAPTIVE FREQUENCY HOPPING IN THE UNLICENSED BANDS

PETAR POPOVSKI, HIROYUKI YOMO, AND RAMJEE PRASAD, AALBORG UNIVERSITY

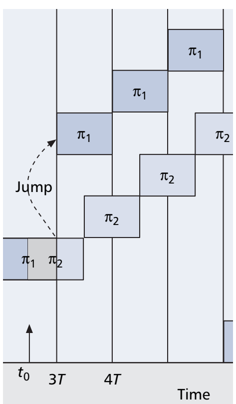

The authors present two novel AFH strategies: Adaptive Frequency Rolling and Dynamic AFH. AFR avoids the self-interference while preserving the dynamics of the spectrum usage. DAFH is a distributed mechanism by which collocated piconets select non-conflicting hopsets, while trying to keep the hopset size as large as possible.

ABSTRACT

Mechanisms based on frequency hopping have been widely used to enable short-range wireless networks to use resources from the unlicensed spectrum without frequency planning. Bluetooth piconet is a prime example of an FH-based network with unlicensed operation. As a price for open access, the piconet may experience adverse interference from other collocated FH piconets or other wireless devices that are transmitting in the same unlicensed band. A basic approach to mitigate this interference is that the piconet applies adaptive FH (AFH) and attempts to hop over a set (hopset) of less interfered channels. On the other hand, the regulation of unlicensed operation sets constraints on possible hopset adaptations. In this article we present two novel AFH strategies: adaptive frequency rolling (AFR) and dynamic AFH (DAFH). AFR avoids self-interference while preserving the dynamics of spectrum usage as required by the current regulation. DAFH is a distributed mechanism by which collocated piconets select nonconflicting hopsets while trying to keep the hopset size as large as possible. DAFH is not completely compliant with current regulations, but the rationale given for its design contains new rules of behavior for the unlicensed spectrum. Both approaches significantly outperform the conventional AFH strategy.

INTRODUCTION

The emerging short-range wireless technologies with ubiquitous usage imply utilization of the unlicensed spectrum. As a price for open access, the unlicensed wireless network may experience adverse interference from collocated wireless devices that are transmitting in the same unlicensed band. A wireless network should exhibit adaptive usage of the unlicensed spectrum so as to attain the best communication performance under the actual interference pattern, while complying with the regulations for unlicensed operation.

Frequency hopping (FH) is a method to enable sharing of the unlicensed spectrum among proximate networks, since it achieves frequency diversity and enables spectrum sharing without frequency planning. FH also decreases interference toward unlicensed devices that operate over a part of the spectrum used by the FH network; this is addressed by defining rules for spectrum usage that limit the radiated power and constrain the occupancy of the channels used in hopping. Bluetooth [1] is a prime example of an FH-based networking technology with unlicensed operation. Bluetooth represents an instance of the wireless personal area network (WPAN), which has been further standardized within the IEEE 802.15 Working Group for WPAN [2]. The FH networks considered in this article inherit their structure from Bluetooth, but the implications of the proposed mechanisms certainly reach beyond Bluetooth. A piconet is a network of devices that share the same FH sequence. A *hopset* is the set of channels used for hopping. Bluetooth uses a hopset of 79 frequencies in the unlicensed industrial, scientific, and medical (ISM) 2.4 GHz band. Recent work has shown that the performance of a piconet can be heavily degraded by interference from collocated piconets or non-FH networks.

Interference in the unlicensed band can be mitigated by using either collaborative or noncollaborative techniques. In collaborative techniques, the interfering entities explicitly exchange data to achieve mutual coordination, while there is no such information exchange in noncollaborative techniques. Adaptive FH (AFH) is an important noncollaborative mechanism considered in [3] to mitigate the interference experienced by an FH-based piconet. An AFH piconet attempts to select a hopset that consists of less interfered channels, while the individual channel occupancy conforms to the regulation [4]. From the AFH perspective, there are three different error sources for a piconet: noise, frequency-static (FS) interference, and frequency-dynamic self-interference. Noise-induced errors are uniform over the channels and cannot be decreased by employing intelligent hopping. FS interference occurs at a group of channels for a time that is considerably longer than the packet duration. A wireless LAN (WLAN) that uses direct sequence spread spectrum (DSSS) is a canonical example of an FS interferer to a Bluetooth piconet [1, 2, 5]. The AFH strategy that combats

1536-1284/06/$20.00 © 2006 IEEE

FS interference is referred to as *standard AFH* and is quite straightforward: The piconet monitors the channel quality in the hopset during a time interval and classifies each channel as good or bad. A bad channel is removed from the hopset for a certain timeout, after which it is brought back to be utilized in the hopset and reassessed. Channel classification schemes [3] may use received signal strength indication (RSSI), packet error rate (PER), and/or carrier sensing; the method that relies on PER is by far the simplest. Frequency-dynamic *self-interference* is intermittent disturbance experienced from a collocated FH-based piconet that has used an identical channel simultaneously [6–8]. The above strategy of "bad channel removal" is not applicable: If a piconet removes a channel where error has been experienced, it is likely that the collocated piconet that caused the error will also remove the same channel. This worsens the throughput, since FH networks self-interfere with smaller hopsets and cause error to each other with higher probability.

In this article we present two novel AFH strategies: adaptive frequency rolling (AFR) [9] and dynamic AFH (DAFH) [10]. Each strategy can be understood as a particular instance of AFH by which the collocated piconets select nonconflicting hopsets in a randomized distributed manner. AFR avoids self-interference while preserving the dynamics of spectrum usage as required by current regulations. The results show that AFR markedly improves the communication performance of collocated piconets. DAFH is significantly different from AFR. DAFH is a distributed mechanism by which collocated piconets select nonconflicting hopsets, while trying to keep the hopset size as large as possible. DAFH is not completely compliant with the current regulation, but the rationale given for its design contains a new rule of behavior for the unlicensed spectrum, which is itself valuable input to future regulations for unlicensed operation.

The article is organized as follows. The next section presents the system model to be used to present AFR and DAFH. We show the simulation results, and the last section concludes the article.

System Model

Our model for an FH-based piconet is based on Bluetooth. The Bluetooth piconet is a star topology with a master and up to seven active slaves. The communication channel is slotted, and the hop selection in each slot is based on a pseudorandom generator determined by the master. The slaves are time- and hop-synchronized to the master. The master uses polling to schedule the packet transmissions such that intrapiconet communication is collision-free, and at each slot only one device transmits.

The ith piconet is denoted π_i, where $i = 1, 2, \ldots$. We define two piconets as *collocated* if the devices in the networks are in each other's transmission range. We model the piconet π_i as a transmitter for which the probability that a packet occurs in a slot is G_i [6]. The piconets are not synchronized with each other such that the slot-starts in different piconets do not coincide. A piconet applies per-packet FH: A frequency is selected at the slot-start and remains constant during the packet transmission. Here we assume that a collision is always *destructive*, resulting in a packet error with probability one if the packet is transmitted in the collocated entities simultaneously at the same frequency [6, 7]. Regarding the channel assessment mechanism, both AFR and DAFH utilize PER measurements. We do not assume any sophisticated mechanism for collision detection — the collided packet is detected by an error detection code. Therefore, the receiver *cannot distinguish* between collision and error due to another error source. Finally, we assume that each piconet utilizes one-slot packets only, which facilitates exposition but is not a necessary condition for the mechanisms to operate.

Our model for FH-based piconet is based on Bluetooth. The Bluetooth piconet is a star topology with a master and up to seven active slaves. The communication channel is slotted and the hop selection in each slot is based on a pseudorandom generator determined by the master.

Adaptive Frequency Rolling

Adaptive frequency rolling [9] is a novel FH method that enables collocated piconets to share channels in an implicit time-division manner. To sketch the basic idea behind AFR, we start with nominal frequency rolling (FR) in which no adaptation is applied. The regulation [4] essentially limits the time for which a channel can be occupied during any interval of duration T_0. Let us assume that by pseudorandom hopping over the complete allocated hopset, such a condition of channel occupancy is satisfied. For the examples in Fig. 1 let us assume that the full set has $F = 80$ channels, numbered 0–79. For both Fig. 1a and 1b, a shaded region denotes the current hopset. During the time a hopset with H channels is used, each channel from the hopset is selected with probability $1/H$. For Fig. 1a the hopset consists of all channels $H = F$, and it is constant over time. For Fig. 1b there are four defined hopsets, each consisting of $H = 20$ frequencies. The hopset changes after $T = T_0/8$; thus, each hopset is used two times during T_0. The average time for which a channel is occupied during T_0 when the whole hopset is used is $T_0/80$. This time is identical for Fig. 1b, since $2 \cdot T/20 = T/10 = T_0/80$. The piconet that does hopping as in Fig. 1b is said to perform FR, since its hopset can be visualized "rolling" in time over the whole set of disposable channels.

FR is an alternative to hopping for achieving identical channel occupancy over a considered timeframe. To introduce AFR, consider piconet π_1 in Fig. 2, and let π_2 become collocated with π_1 at instant t_0. Both piconets have identical T and are rolling at the same pace, with a hopset of size $H = 10$. On the other hand, the piconets are asynchronous in a sense that the instants of hopset change in the two piconets do not coincide. Piconets π_1 and π_2 detect conflict through experiencing excessive PER during the time from t_0 to $3T$. Let us assume that, having detected excessive PER, π_1 decides to interrupt the nominal FR and changes the hopset by a random jump. On the other hand, π_2 continues with nominal rolling. It can be seen from Fig. 2 that after there is an initial conflict and π_1 adapts the rolling via a random jump, self-interference is completely eliminated.

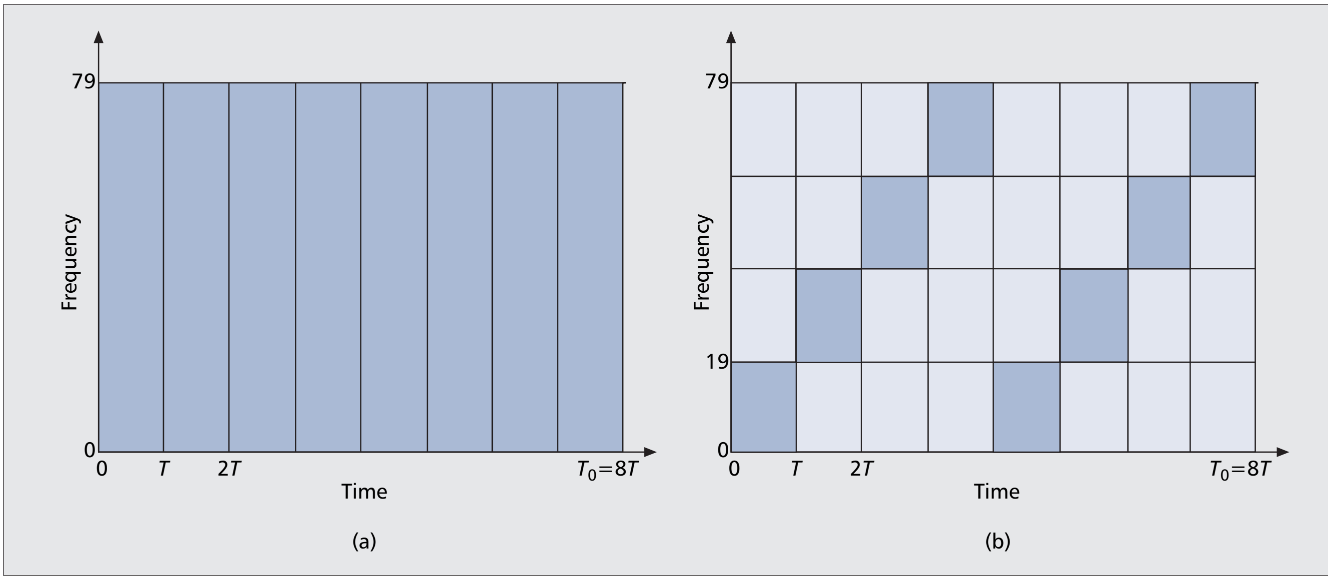

■ **Figure 1.** *a) Pseudorandom FH over the full hopset of 80 frequencies; b) frequency rolling with a hopset of 20 frequencies.*

AFR is designed as a hopping method that is compliant with the regulations for unlicensed operation. Let us first define the channel occupancy within some interval T_0 as the total time within T_0 for which the piconet uses that channel. In essence, the regulation for unlicensed operation puts upper limits on individual channel occupancy. Considering that we have taken the reference FH technology to be Bluetooth, we focus on the regulation in the ISM band. Here is the regulation statement for the FH systems in the unlicensed ISM band [4]:

(**RS1**) *Frequency hopping systems in the 2400–2483.5 MHz band shall use at least 15 nonoverlapping channels. The average time of occupancy on any channel shall not be greater than 0.4 seconds within a period of 0.4 seconds multiplied by the number of hopping channels employed. Frequency hopping systems which use fewer than 75 hopping frequencies may employ intelligent hopping techniques to avoid interference to other transmissions. Frequency hopping systems may avoid or suppress transmissions on a particular hopping frequency provided that a minimum of 15 non-overlapping channels are used.*

A straightforward way to be in agreement with RS1 is to apply pseudorandom FH (PFH) over all F disposable channels. Another rather obvious way to achieve compliance with RS1 is that each piconet hops pseudorandomly over a set of at least 15 channels which is noninterfered. This permits a situation to be achieved in which at most $\lfloor 79/15 \rfloor = 5$ collocated piconets have nonintersecting hopsets and thereby avoid self-interference. The goal of AFR is to accommodate more than five noninterfering WPANs, while keeping the channel occupancy lower or equal to that permitted by RS1. The maximal occupancy caused by a piconet compliant with RS1 occurs when the hopset has only 15 channels. Hence, in the AFR design we have adopted the following statement:

(**RS2**) *A frequency hopping system should use the disposable channels in such a way that, in any interval of 6 seconds, no channel is used for more than 0.4 seconds.*

It can be proven [9] that a system designed according to RS2 cannot produce channel occupancy higher than that produced by a system with 15 channels designed according to RS1.

Design of the AFR Parameters

We first discuss nominal FR in the absence of FS interference. The set of all F frequency channels over which rolling is applied is called a *rollset*. The *roll interval* has a duration of T slots during which the piconet does not change its hopset. The hopset within a given roll interval is defined by a subset of H adjacent channels from the rollset. The size H is taken to be constant for all piconets, but this is not a necessary condition. At each slot, the piconet chooses pseudorandomly and uniformly a channel from its current hopset. After the termination of the roll interval, the piconet changes to a new hopset of size H by shifting (modulo F) each channel from the old hopset for a value of Δr, called *roll step*, which is constant and identical for all piconets. For example, if $H = 3$, $F = 10$, $\Delta r = 2$, and the current hopset is given by {6, 7, 8}, the next hopset is {8, 9, 0}. The parameters T and Δr are predefined and necessarily constant across all piconets to ensure implicit time division. To generate the FR hopping sequence, the piconet master shares with all the slaves a seed to generate pseudorandom numbers from the set $\{0, 1, \ldots H - 1\}$, which uniquely map the numbers to the channels from the current hopset. Since the roll step is predefined, at the end of the current roll interval there is no need to explicitly convey the offset of the next hopset to the slaves.

If the PER during rolling interval exceeds a threshold, the piconet master is triggered to interrupt the nominal FR and prematurely change the hopset, which is in fact AFR. Upon such a change, a *random jump* is applied, and the offset of the hopset is changed in a randomized way. The offset value and timing associated with the jump should be explicitly disseminated to the slaves. The dissemination should be done reliably by either unicasting to each slave (used for standard AFH in Bluetooth) or broadcasting to all slaves [9]. Note that, in absence of a random jump, the hopset change (timing and offset) due to the nominal rolling is known to all piconet members in a predefined manner, such that no information needs to be conveyed within the piconet, and there is no overhead associated with such a hopset change. The PER threshold used for triggering should be chosen high enough to minimize false triggering from noise-induced errors. Furthermore, we have introduced randomization in the threshold choice in order to avoid symmetric behavior of different piconets.

AFR should be designed in a way which ensures that overall piconet behavior is statistically compliant with the channel occupancy limit set by RS2. This includes defining timing constraints for piconets as well as restrictions on the selection of the random jump [9]. These constraints in hopset adaptation may impair the overall performance of the AFR, as can be seen later. There are trade-offs in selecting hopset size H and roll interval T. A smaller value of H enables coexistence of more noninterfering piconets, but leads to low *short-term frequency diversity* and thereby possible severe degradation during the rolling interval if the channels from the hopset are interfered with or there is frequency-selective fading. While $H = 1$ is acceptable from a regulatory viewpoint, using $H = 1$ may lead to a deadlock in the operation of AFR. This is because prior to a random jump, the master needs to send a certain amount of packets correctly to the slaves in order to inform them about the upcoming random jump. If $H = 1$ and two collocated piconets are using the same channel, no information can be sent within each of the piconets unless using some ALOHA-like mechanism to resolve the contention. On the other hand, if $H > 1$ the pseudorandom frequency selection inherently avoids the deadlock in contention. The choice of roll interval T should be made to ensure compliance with RS2, but it should also allow reliable PER estimation and reliable dissemination of hopset change upon triggering to be achieved. An issue that could be raised is the relative clock drift among the piconets. A piconet with a faster clock can "reach" the piconet with a slower clock and again have a conflicting hopset. AFR is robust to such an occurrence, since it will simply again apply a random jump adaptation.

Finally, we outline the modification of AFR to cope with FS interference. With AFR the piconets implicitly cooperate to avoid interference, but the FS interferer is not assumed to cooperate. Thus, the only applicable strategy for the FS interferer is to change the rollset by "bad channel removal." We also point out an additional necessary modification of the AFR: The channels interfered with by an FS interferer are usually contiguous, and if they cause interference to the piconet at all the H channels, the master has a problem telling the slaves the randomized hopset change. To cope with this problem as well as introduce frequency diversity during the roll interval, the rolling is done over the interleaved set of channels where the interleaving is identical for all piconets and is done according to some predefined permutation.

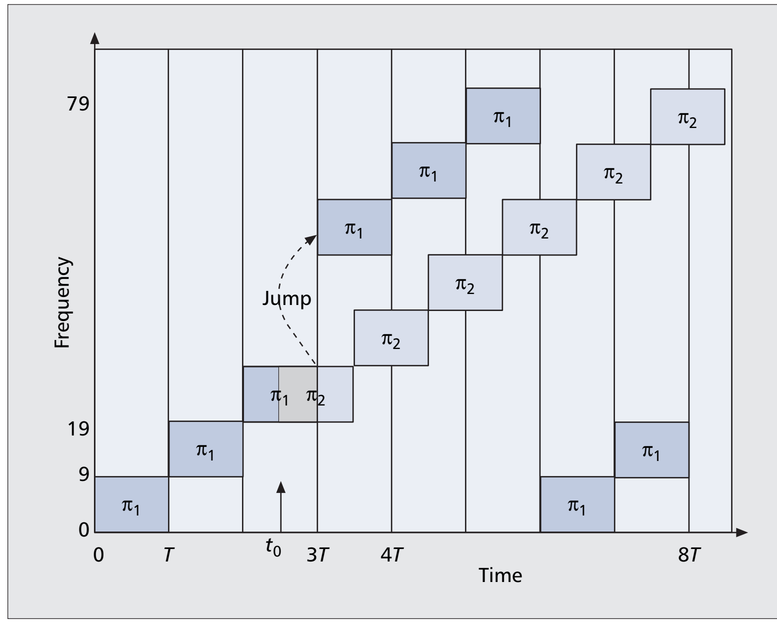

■ **Figure 2.** *Frequency rolling for two collocated piconets π_1 and π_2. The time reference* T, 2T … *is given for π_1. At the beginning of* 3T, *piconet π_1 adapts the rolling through a random jump.*

In order to enable rollset change, we introduce AFR with probing (AFR-P). In the *rolling state*, the piconet performs AFR over the current rollset, while in the *probing state*, the piconet applies pseudorandom FH (PFH) over the whole set of F channels. While probing, the master attempts to detect which channels are experiencing FS interference, remove them, and perform over AFR a reduced rollset. Rolling is the stable state of the piconet, while probing is transient and lasts for a fixed duration of T_{probe} slots. Too short T_{probe} leads to unreliable PER statistics for each channel, but too long T_{probe} brings prolonged periods of self-interference, thus wiping out the benefits of the AFR mechanism. When the piconet is triggered in the rolling state, the master instructs the slave either to do a random jump within the current rollset or that the piconet should start probing. After a T_{probe} in the probing state, the master informs the slaves about the new rollset as well as the first hopset to be used after returning to the rolling state. For running AFR-P the piconet can have an additional FH generator used for the probing state, and both FH generators should be updated in each slot. Compared to standard AFH, the the AFR-P algorithm does not include an explicit mechanism for recovery of the removed channels, since it recovers the channels by being triggered to enter into the probing state.

DYNAMIC ADAPTIVE FREQUENCY HOPPING

The bad channel removal strategy induced by an FS interferer seems unfair to the piconet, and it is natural to ask why the piconet should be

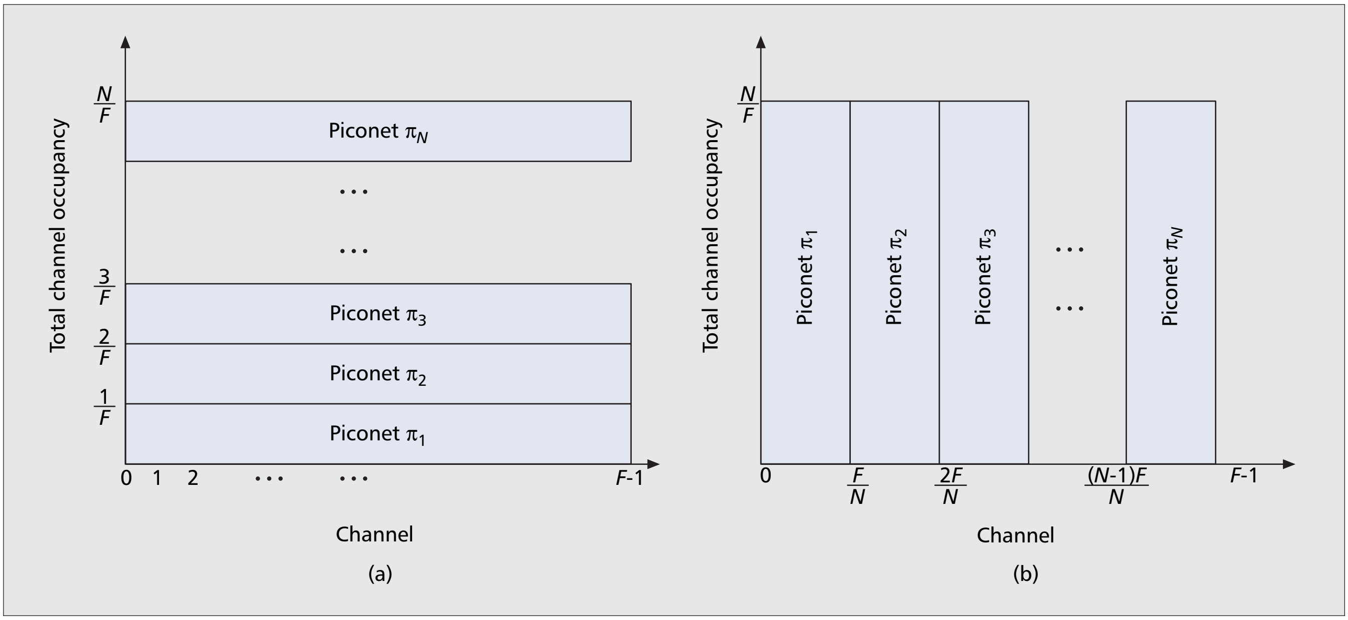

Figure 3. *Total occupancy of frequency channels, estimated as a fraction of time a channel is occupied: a) when each piconet hops over all* F *channels; b) when the piconets use* N *orthogonal hopsets, and each hopset contains* F/N *channels.*

"polite" according to the regulation if there is a noncooperative collocated entity that compels the piconet to temporarily remove some channels from its hopset. In essence, the regulation is created according to some *etiquette* [11] by which the FH system should utilize the unlicensed spectrum. This motivates us to think in the direction of formulating alternative etiquette rules. Such rules should still prevent unfair opportunistic behavior in the unlicensed spectrum, but allow more flexibility in defining the algorithms for resource utilization.

To illustrate how such an etiquette rule can be introduced, we consider two piconets π_1 and π_2 that are performing PFH over the whole set **F** of F channels. They are self-interfering, while the fraction of time for which a channel i is occupied (which we call fractional occupancy) is

$$1-\left(1-\frac{1}{F}\right)^2=\frac{2}{F}-\frac{1}{F^2}\approx\frac{2}{F}. \quad (1)$$

Now let the hopset of π_1 be only half of the set of all available channels, and let the hopset of π_2 consist of the other $F/2$ channels from **F** that are not used by π_1. The fractional occupancy of each channel i is $2/F$, which is practically equal to the occupancy from Eq. 1. Nevertheless, in this case the piconets are not self-interfering. This simple example shows that it is possible to have a set of noninterfering collocated piconets, while not changing the interference to the environment caused from the set of piconets observed as a collective entity. The same observation with N piconets is illustrated in Fig. 3. This is the essence of the etiquette rule we introduce below.

Etiquette rule for DAFH: *Let* N *collocated entities apply adaptive frequency hopping to avoid mutual interference. Let each entity have the same traffic load. Then the occupancy of each individual frequency with adaptive hopping should remain as close as possible to the occupancy induced by the same* N *collocated entities when they do not apply any mechanism to avoid mutual interference*.

The proposed DAFH algorithms exhibit behavior that follows this etiquette rule rather than the current regulation, which has been the foundation for the AFR mechanisms. The collocated piconets that apply DAFH aim to select hopsets to avoid self-interference, while making their best effort to keep the interference equivalent to that caused by the same set of piconets if performing PFH. Since the DAFH contains the strategy of bad channel removal, it offers inherent immunity to an FS interferer.

The basic strategy applied in DAFH is a binary search for a hopset that offers smaller PER. The set of admissible hopsets depends on the total number of available channels and the maximal level of binary divisions. For example, if $F = 8$ and the maximal level $L = 2$, the set of admissible hopsets is given in Table 1. The target behavior of DAFH is achieved by a combination of two mechanisms: *hopset reduction* and *hopset doubling*.

Each piconet starts to operate using the full hopset **F** (i.e., the admissible hopset at level 0). If the PER exceeds the threshold, the piconet is triggered and reduces its hopset. The idea of reduction is that the triggered piconet should reduce its hopset by randomly selecting the new hopset to be either the left or right half of the current hopset. If the piconet is again triggered, it reduces (halves) the hopset further. Hence, if two piconets select the same hopset, they resolve the conflict by splitting the hopset and randomly selecting half of the split hopset. Such an operation relies on the same principles utilized in the classical collision resolution algorithms [12]. Splitting cannot continue indefinitely, and if the piconet is triggered while using a hopset at the maximal level L, it randomly selects another hopset from level L. Hopset doubling is complementary to reduction. If the piconet operates

Lev el	No. hopsets in the level	Channels in the hopsets			
0	1	0 1 2 3 4 5 6 7			
1	2	0 1 2 3		4 5 6 7	
2	4	0 1	2 3	4 5	6 7

■ **Table 1.** *Set of admissible hopsets for* F = 8 *and* L = 2.

with a hopset at level $l > 0$ and is not being triggered for a time T_D, the piconet doubles its hopset by randomly selecting a hopset from level $l - 1$. If the doubling period T_D is too short, the effect of orthogonalization by reduction will be hindered, and a large overhead will be introduced. If T_D is too long, the interferers that forced the piconet to reduce the hopping may not be collocated anymore, such that the piconet induces unnecessarily large occupancy at the channels of the reduced hopset. Note that randomization in reduction/doubling ensures fairness among the piconets over a longer period.

It is highly likely that DAFH does not violate the current regulations when interference to the piconet is caused by an FS interferer. When there are multiple collocated piconets, it may seem that DAFH violates the regulation by using DAFH, since a piconet may use a hopset smaller than the smallest allowed by the Federal Communications Commission (FCC). However, if the collocated piconets are regarded as a collective entity, DAFH employs a best effort to minimize change of interference pattern.

To generate the DAFH hopping sequence, the master shares with the slaves the seed to generate a pseudorandom sequence which consists of the integers that are uniformly chosen from the full set. The dissemination of hopset change information can be done in identical ways as for AFR, but it is even easier in DAFH due to the absence of time constraints. When a hopset is subset of the full set, the hopping sequence is obtained by a predefined many-to-one mapping. Two key design issues in hopset reduction are the triggering threshold and hopset selection. We have introduced two threshold variants of DAFH [10]: constant threshold (DAFH-CT) and adaptive threshold (DAFH-AT). In DAFH-CT the PER threshold is always kept above the estimated PER of noise-induced errors. In DAFH-AT the threshold increases as the hopset becomes smaller (i.e., the level of the hopset increases). The motivation behind the adaptive threshold is to prevent unnecessary hopset reduction: if a piconet interferes with a piconet with a larger hopset, the latter is more likely to reduce is hopset. Hopset selection upon reduction or doubling can be made to be either uniformly randomized or randomized in a way that is biased by the gathered channel statistics. For example, if the piconet has assessed that PER within the left half of the current subset is significantly higher, it selects the right half with higher probability.

Simulation Results

In the simulation scenario the piconets arrive at a hotspot (e.g., lounge at airport, conference room) and leave the hotspot after a random dwell time. The piconet arrivals form a Poisson process of rate λ. The dwell time of the piconet in the hotspot is shifted exponential distribution with the average of $\overline{T}$ [s]. We set the minimal dwell time to be 20 [s]. In this case, the average number of piconets at the hotspot, denoted $\overline{N}$, can be obtained as $\overline{N} = \lambda \times (\overline{T} + 20)$. In the simulation, we fix $\overline{T}$ at 60 [s], and change the average number of piconets by setting different value of λ. In order to challenge the operation of AFR and DAFH, we intentionally choose the values for the dwell time to be short, which creates the worst-case scenarios for the AFR/DAFH operation. If the dwell time values are larger, the piconets have enough time to separate their hopsets and operate for a considerable period without self-interference and triggering. The full set contains 79 channels, as in Bluetooth. For DAFH, the maximal level of splitting is $L = 4$; the admissible hopsets are created as if the number of channels is $F = 80$, and then the 80th channel is removed from the hopsets that contain it. A packet fully covers the slot, and the piconets are fully loaded at $G = 1$. The reference system in which the piconet selects the frequency pseudorandomly from the full channel set is denoted PFH. The key performance mea-

In general, both the AFR and the DAFH mechanism can be understood as ingredients for designing a hopping scheme that enables self-organized division of the resources among the interfering systems. Such mechanisms will gain more importance .

Pseudorandom frequency hopping (PFH)	The hopset is kept constant, and a channel is chosen pseudorandomly from the hopset.
Adaptive frequency hopping (AFH)	The hopset is changed by removing the channels that experience excessive error.
Adaptive frequency rolling (AFR)	Usage of a small hopset that, in absence of errors, is regularly changed after a short predefined time. Upon excessive errors, the hopset is changed in a randomized manner.
Adaptive frequency rolling with probing (AFR-P)	Extension of AFR that deals with FS interference by suspending the use of "bad" channels for a longer time period.
Dynamic adaptive frequency hopping (DAFH)	Upon interference, a randomized binary splitting of the hopset is applied in order to avoid conflict with the interfering piconet.

■ **Table 2.** *Summary of the different hopping methods.*

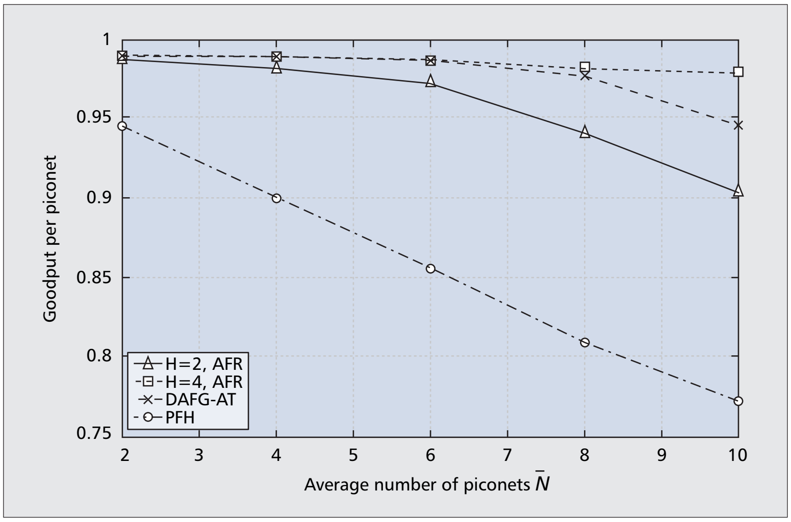

■ **Figure 4.** *Comparison of average goodput of DAFH and AFR with PER due to noise errors of* p_n *= 1 percent.*

sure is the goodput, which is defined as the fraction of time used by successful packet transmissions excluding overhead packets to convey the information of the hopset change. The average goodput of a piconet is obtained by averaging the goodput of all the piconets that have been in the hotspot during the simulation.

Figure 4 compares, in the absence of FS interference, the goodput of DAFH-AT, AFR, and conventional PFH over the full channel set. The PER due to channel noise is 1 percent. Hopset sizes for AFR are $H = 2$ and $H = 4$. First, it can be seen that both AFR and DAFH significantly outperform piconet operation with PFH. AFR with $H = 2$ always outperforms AFR with $H = 4$ and DAFH-AT. However, recall that $H = 2$ may lead to severe temporary goodput degradation due to low short-term frequency diversity. The average goodput performance of DAFH-AT always outperforms the AFR variant with $H = 4$. This is because DAFH-AT does not use time constraints to satisfy the regulation, and in case of increased PER it promptly attempts to change the hopset.

To evaluate the impact of an FS interferer at the hotspot, we have included standard AFH as a reference. The set of channels influenced by the FS interferer are selected to simulate the interference of IEEE 802.11b with Bluetooth. Whenever a piconet chooses to transmit on a channel with FS interference, the probability that the packet is erroneously received is equal to the activity factor A_s of the FS interferer, and we use $A_s = 0.7$. Figure 5 compares the average goodput per piconet for systems with DAFH-AT, AFR-P with $H = 2$ and $H = 3$, standard AFH (denoted AFH), and PFH. DAFH-AT offers the best goodput performance. Nevertheless, AFR-P appears inferior to DAFH-AT since it applies stringent constraints in hopset adaptation that yield compliance with the regulation for unlicensed FH. Even more, DAFH-AT as considered here is simpler to implement than standard AFH or AFR-P, since it does not use relative PER for each frequency channel but only two variables to monitor the state in the left and right halves of the current hopset. It can be seen that the gain of DAFH over standard AFH is higher than the gain of AFH over the case when no AFH is applied for practical values of the average number of collocated piconets $\bar{N} <$ 10. Comparing Figs. 4 and 5, we can see that the presence of an FS interferer leads to larger goodput degradation for DAFH/AFR as the average number of collocated piconets increase. This is because the FS interferer compels the piconets to resolve the conflicts within a smaller subset of channels, which degrades the success of conflict resolution.

Conclusion

In this article we have presented two novel strategies for adaptive frequency hopping (AFH) applied by the piconets that operate in the unlicensed band: adaptive frequency rolling (AFR) and dynamic AFH (DAFH). AFR complies with the current regulation for FH operation in the unlicensed spectrum. An important novelty brought by AFR is the design methodology for producing FH patterns that are provably compliant with the requirements for FH in the unlicensed ISM band. We have also discussed how the basic AFR approach should be modified to combat FS interference and presented AFR with probing (AFR-P). The main problem of AFR is that the better it is at avoiding self-interference, the smaller the short-term frequency diversity. A fundamental difference between AFR and DAFH occurs at the basic design premises: While AFR has been designed to comply with the current regulation for unlicensed operation, DAFH relies on arguments that can affect future regulations. The second difference between DAFH and AFR is that in DAFH the piconet always attempts to pseudorandomly select the frequency from a hopset that is as large as possible, while in AFR the pseudorandom selection is always done from a hopset with small size H. This means that DAFH makes a best effort to maximize the short-term frequency diversity. Finally, the DAFH mechanism can inherently combat FS interference.

In general, both the AFR and DAFH mechanisms can be understood as ingredients for designing a hopping scheme that enables self-organized division of the resources among the interfering systems. Such mechanisms will gain more importance as there is increasing crowding of the unlicensed bands without frequency preplanning.

References

[1] Bluetooth Std. Spec., Rev. 1.2, Nov. 2003.

[2] IEEE 802.15 WG on Wireless Personal Area Networks, Sept. 2004; http://www.ieee802.org/15/

[3] IEEE Std. 802.15.2, "Coexistence of Wireless Personal Area Networks with Other Wireless Devices Operating in Unlicensed Frequency Band," 2003; http://standards.ieee.org/

[4] FCC, "Operation within the Bands 902–928 MHz, 2400–2483.5 MHz, and 5725–5850 MHz," Pt. 15: Radio Frequency Devices, Oct. 2002.

[5] C. F. Chiasserini and R. Rao, "Coexistence Mechanisms for Interference Mitigation in the 2.4-GHz ISM Band," *IEEE Trans. Wireless Commun.*, vol. 2, Sept. 2003, pp. 964–75.

[6] A. El-Hoiydi, "Interference between Bluetooth Networks — Upper Bound on the Packet Error Rate," *IEEE Commun. Lett.*, vol. 5, no. 6, June 2001, pp. 245–47.

[7] T.-Y. Lin and Y.-C. Tseng, "Collision Analysis for a Multi-Bluetooth Picocells Environment," *IEEE Commun. Lett.*, vol. 7, no. 10, Oct. 2003, pp. 475–77.
[8] F. Florén *et al.*, "Throughput of Strongly Interfering Slow Frequency-Hopping Networks," *IEEE Trans. Commun.*, vol. 52, no. 7, July 2004.
[9] P. Popovski *et al.*, "Frequency Rolling: A Cooperative Frequency Hopping for Mutually Interfering WPANs," *Proc. ACM Mobihoc '04*, Tokyo, Japan, May 2004, pp. 199–209.
[10] P. Popovski, H. Yomo, and R. Prasad, "Dynamic Adaptive Frequency Hopping for Mutually Interfering Wireless Personal Area Networks," *IEEE Trans. Mobile Comp.*, to appear.
[11] J. M. Peha, "Wireless Communications and Coexistence for Smart Environments," *IEEE Pers. Commun.*, vol. 7, no. 5, Oct. 2000, pp. 66–68.
[12] J. L. Massey, *Collision-Resolution Algorithms and Random-Access Communications*, CISM Courses and Lectures, Springer-Verlag, 1981, no. 265, pp. 73–137.

Biographies

PETAR POPOVSKI [M] (petarp@kom.auc.dk) received a Dipl.-Ing. in electrical engineering and an M.Sc. in communication engineering from the Faculty of Electrical Engineering, Sts. Cyril and Methodius University, Skopje, Macedonia, in 1997 and 2000, respectively. He received a Ph.D. degree from Aalborg University, Denmark, in 2004. From 1998 to 2001 he was a teaching and research assistant at the Institute of Telecommunications, Faculty of Electrical Engineering in Skopje. He is currently an assistant professor with the Department of Communication Technology, Aalborg University. His research interests are related to the PHY-MAC aspects of wireless protocols, wireless sensor networks, random access protocols, and network coding.

HIROYUKI YOMO [M] received a B.S. degree in communication engineering from the Department of Communication Engineering, Osaka University, Japan, in 1997, and M.S. and Ph.D. degrees in communication engineering from the Department of Electronic, Information, and Energy Engineering, Graduate School of Engineering, Osaka University, in 1999 and 2002, respectively. From April 2002 to March 2004 he was a postdoctoral fellow in the Department of Communication Technology, Aalborg University. From April to September 2004 he was at Internet System Laboratory, NEC Corporation, Japan. Since October 2004 he has been an assistant research professor with the Center for TeleInfrastructure (CTIF), Aalborg University. His main research interests are access technologies, radio resource management, and link-layer techniques in the area of short-range communication, cellular networks, cognitive radio, and sensor networks.

RAMJEE PRASAD [SM] is a distinguished educator and researcher in the field of wireless information and multimedia communications. From February 1988 to May 1999 he was with the Telecommunications and Traffic-Control Systems Group of Delft University of Technology (DUT), The Netherlands, where he was actively involved in the area of wireless personal and multimedia communications (WPMC). He was head of the Transmission Research Section of International Research Centre for Telecommunications Transmission and Radar (IRCTR) and also Founding Program Director of the Centre for Wireless Personal Communications (CWPC). In June 1999 he became Wireless Information Multimedia Communications chair and co-director of the Centre for PersonKommunikation at Aalborg University. In January 2004 he was the Founding Director of the Centre for Teleinfrastruktur (CTIF). He has published over 500 technical papers, and authored and co-edited 15 books about wireless multimedia communications (Artech House). His research interest lies in wireless networks, packet communications, multiple access protocols, adaptive equalizers, spread-spectrum CDMA systems, and multimedia communications. He is the founding chairperson of the European Centre of Excellence in Telecommunications known as the HERMES Partnership. He is a fellow of the IEE, a fellow of IETE, a member of NERG, and a member of the Danish Engineering Society (IDA). He is an advisor to several multinational companies.

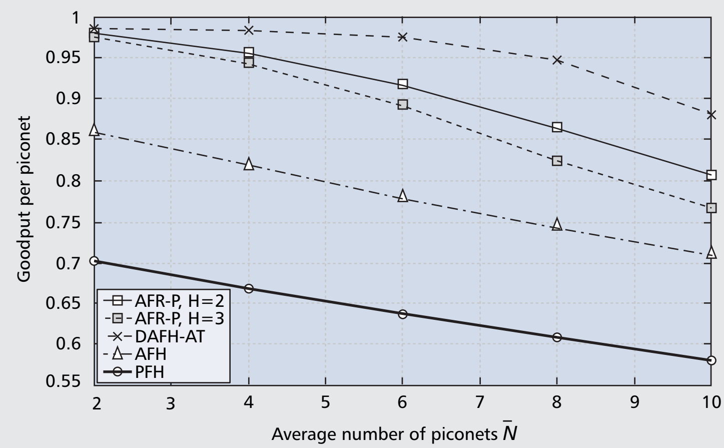

Figure 5. *Average goodput against average number of piconets with an FS interferer that covers 22 frequencies and has activity $A_s = 0.7$. The piconets have G = 1, and PER due to noise errors is 1 percent.*